DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1. (Currently Amended) A server system comprising:

a high performance spinning-type hard drive for storing data; and

a power controlled spinning-type hard drive for storing data, wherein an entirety of the power controlled spinning-type hard drive is configured for lower power consumption as compared to the high performance spinning-type hard drive; and

a policy manager that assesses multiple metadata analysis of a data set to be stored into one of the high performance hard drive or the power controlled hard drive, the multiple metadata analysis including all of; (a) an analysis of usage rate metadata of the data set to be stored, (b) an analysis of content type metadata of the data set to be stored, and (c) an analysis of owner and consumer metadata of the data set to be stored, wherein the policy manager receives the multiple metadata analysis of (a), (b), and (c) and makes a determination whether the data set is to be stored in the high performance hard drive or in the low power hard drive, the determination based on rules to eliminate contention between the multiple metadata analysis, the policy manager implementing rules of an application policy and a file system policy, the application policy defining power management directives within an application, and the file system policy including rules associated with a file system;

wherein the data set to be stored is placed on the power controlled spinning-type hard drive after classification of the data into the characteristics of a fully distinguished file name, creator, owner, consumer, distribution lists, access control list detail, certificates, signature attributes, protocols, content resolution, encoding technique, encryption technique, key properties, internal subjects, keywords, content tags, assemblies, associations to other files, replication, caching, directory and related database extensible properties.

2. (Original) The server system of claim 1, wherein the power controlled spinning–type hard drive achieves a lower power consumption by spinning up the power controlled drive

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

upon receipt of a service request and spinning down the power controlled drive after providing service.

3-6.   (Cancelled)

7.   (Previously Presented) The server system of claim 1, wherein the data set to be stored comprises audio or video files.

8.   (Cancelled)

9.   (Original) The server system of claim 1, wherein the high performance drive spins continuously.

10.   (Previously Presented) The server system of claim 1, wherein a subset of the data set is moved to power controlled drive in accordance with a predetermined condition, the data set moved after being stored in the high performance hard drive.

11.   (Original) The server system of claim 10, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, the number of times the subset of data has been requested in a period of time, the data size, users of the data, a file name, the data types, and internal content.

12.   (Previously Presented) The server system of claim 1, wherein a subset of the data set to be stored is moved to the high performance drive in accordance with a predetermined condition, the data set moved after being stored in the power controlled hard drive.

13.   (Original) The server system of claim 12, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

14.   (Currently Amended) A method for managing data on a server, the method comprising:

receiving a data set to be stored, the data set to be stored into one of a high performance hard drive or a power controlled hard drive;

performing multiple metadata analysis of the data set, the multiple metadata analysis including all of; (a) an analysis of usage rate metadata of the data set to be stored, (b) an analysis of content type metadata of the data set to be stored, and (c) an analysis of owner and consumer metadata of the data set to be stored;

receiving the multiple metadata analysis of (a), (b), and (c) by a policy manager; and

determining, using the policy manager, whether the data set is to be stored in the high performance hard drive or in the low power hard drive, the determination based on rules to eliminate contention between the multiple metadata analysis, the policy manager implementing rules of an application policy and a file system policy, the application policy defining power management directives within an application, and the file system policy including rules associated with a file system;

wherein the step of performing multiple metadata analysis further comprises analyzing the data set according to the characteristics of a fully distinguished file name, creator, owner, consumer, distribution lists, access control list detail, certificates, signature attributes, protocols, content resolution, encoding technique, encryption technique, key properties, internal subjects, keywords, content tags, assemblies, associations to other files, replication, caching, directory and related database extensible properties.

15.   (Cancelled)

16.   (Previously Presented) The method of claim 14, wherein the step of performing the analysis of usage rate metadata comprises:

determining a storage location for the data set by comparing the request rate of the data set to a predetermined threshold;

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

assigning data that has a request rate above the predetermined threshold to the high performance hard drive; and

assigning data that has a request rate below the predetermined threshold to the power controlled hard drive.

17.   (Previously Presented) The method of claim 14, wherein the step of performing the analysis of usage rate metadata comprises:

determining a storage location for the data set by comparing the last request date of data to a predetermined date;

assigning data that has a last request date after the predetermined date to the high performance hard drive; and

assigning data that has a last request date before the predetermined date to the power controlled hard drive.

18.   (Previously Presented) The method of claim 14, wherein the step of performing the analysis of usage rate metadata comprises:

determining a storage location for the data set by comparing the creation date of data to a predetermined date;

assigning data that has a creation date after the predetermined date to the high performance hard drive; and

assigning data that has a creation date before the predetermined date to the power controlled hard drive.

19.   (Previously Presented) The method of claim 14, wherein the data set comprises audio or video.

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

20.   (Previously Presented) The method of claim 14, wherein the data set comprises documents.

21.   (Original) The method of claim 14, further comprising spinning the high performance drive continuously, and maintaining the power controlled drive in a reduced performance state unless actively accessing data stored thereon.

22.   (Previously Presented) The method of claim 14, further comprising moving a subset of the data set to the power controlled drive in accordance with a predetermined condition, the data set moved after being stored in the high performance hard drive.

23.   (Original) The method of claim 22, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

24.   (Previously Presented) The method of claim 14, further comprising moving a subset of the data set to the high performance drive in accordance with a predetermined condition.

25.   (Original) The method of claim 24, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

26.   (Currently Amended) A computer-readable storage medium having computer-executable instructions for performing a method of managing data on a server, the method comprising:

receiving a data set to be stored, the data set to be stored into one of a high performance hard drive or a power controlled hard drive;

performing multiple metadata analysis of the data set, the multiple metadata analysis including all of (a) an analysis of usage rate metadata of the data set to be stored, (b) an

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

analysis of content type metadata of the data set to be stored, and (c) an analysis of owner and consumer metadata of the data set to be stored;

receiving the multiple metadata analysis of (a), (b), and (c) by a policy manager; and

determining, using the policy manager, whether the data set is to be stored in the high performance hard drive or in the low power hard drive, the determination based on rules to eliminate contention between the multiple metadata analysis, the policy manager implementing rules of an application policy and a file system policy, the application policy defining power management directives within an application, and the file system policy including rules associated with a file system;

wherein the step of performing multiple metadata analysis further comprises analyzing the data set according to one or more of the characteristics of a fully distinguished file name, creator, owner, consumer, distribution lists, access control list detail, certificates, signature attributes, protocols, content resolution, encoding technique, encryption technique, key properties, internal subjects, keywords, content tags, assemblies, associations to other files, replication, caching, directory and related database extensible properties.

27.   (Cancelled)

28.   (Currently Amended) The computer-readable storage medium of claim 26, further comprising the step of spinning the high performance drive continuously, and maintaining the power controlled drive in a reduced performance state unless actively accessing data stored thereon.

29.   (Currently Amended) The computer-readable storage medium of claim 26, further comprising the step of moving a subset of the data set to the power controlled drive in accordance with a predetermined condition, the data set moved after being stored in the high performance hard drive.

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

30.   (Currently Amended) The computer-readable storage medium of claim 29, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

31.   (Currently Amended) The computer-readable storage medium of claim 26, further comprising the step of moving a subset of the data set to the high performance drive in accordance with a predetermined condition, the data set being moved after being stored in the power controlled hard drive.

32.   (Currently Amended) The computer-readable storage medium of claim 31, wherein the predetermined condition is based on at least one of the last access date of the subset of data, the current date, and the number of times the subset of data has been requested in a period of time.

33.   (Original) The system of claim 1, wherein a database is divided such that a first portion of the database, having a first set of attributes, is stored in the high performance hard drive, and a second portion of the database, having a second set of attributes, is stored in the power controlled hard drive.

34.   (Previously Presented) The system of claim 33, wherein the first set of attributes comprises a first set of meta-directives and the second set of attributes comprises a second set of meta-directives.

35.   (Original) The system of claim 1, wherein the high performance hard drive comprises one of a remote storage area network (SAN) device and a network attached storage (NAS) device, wherein power management directives control the storage of data across multiple storage devices.

36.   (Original) The system of claim 1, wherein the power controlled hard drive comprises one of a remote storage area network (SAN) device and a network attached storage (NAS)

DOCKET NO.: MSFT-2933/306027.01
Application No.: 10/823,918
Office Action Dated: November 14, 2007

PATENT
REPLY FILED UNDER EXPEDITED
PROCEDURE PURSUANT TO
37 CFR § 1.116

device, wherein power management directives control the storage of data across multiple storage devices.

37.   (Original) The system of claim 1, wherein the power controlled hard drive comprises one of a remote storage area network (SAN) device and a network attached storage (NAS) device, wherein power management directives are used to control power management events in storage devices.

38.   (Original) The system of claim 1, wherein the power controlled hard drive is a redundant array of independent disks (RAID).

39.   (Previously Presented) The system of claim 38, wherein the second set of data in the power controlled hard drive is distributed across the redundant array using a stripping algorithm.

40.   (Original) The system of claim 39, wherein the stripping algorithm reduces the power up impact to a smaller set of drives upon data access operations.